

# A Framework for Selecting a Selection Procedure

ROLF WAEBER, PETER I. FRAZIER, and SHANE G. HENDERSON, Cornell University

For many discrete simulation optimization applications, it is often difficult to decide which Ranking and Selection (R&S) procedure to use. To efficiently compare R&S procedures, we present a three-layer performance evaluation process. We show that the two most popular performance formulations, namely the Bayesian formulation and the indifference zone formulation, have a common representation analogous to convex risk measures used in mathematical finance. We then specify how a decision maker can impose a performance requirement on R&S procedures that is more adequate for her risk attitude than the indifference zone or the Bayesian performance requirements. Such a performance requirement partitions the space of R&S procedures into acceptable and nonacceptable procedures. The minimal computational budget required for a procedure to become acceptable introduces an easy-to-interpret preference order on the set of R&S policies. We demonstrate with a numerical example how the introduced framework can be used to guide the choice of selection procedure in practice.

Categories and Subject Descriptors: I.6.1 [Simulation and Modeling]: Simulation Theory; G.4 [Mathematics of Computing]: Mathematical Software—*Algorithm design and analysis*

General Terms: Performance, Algorithms, Theory

Additional Key Words and Phrases: Ranking and selection, robustness analysis, decision theory, indifference zone, Bayesian, convex risk measures, acceptable procedures

## ACM Reference Format:

Waeber, R., Frazier, P. I., and Henderson, S. G. 2012. A framework for selecting a selection procedure. *ACM Trans. Model. Comput. Simul.* 22, 3, Article 16 (August 2012), 23 pages.  
DOI = 10.1145/2331140.2331144 <http://doi.acm.org/10.1145/2331140.2331144>

## 1. INTRODUCTION

When stochastic simulation is used to identify the best option among several alternatives, Ranking and Selection (R&S) procedures are often implemented to allocate the simulation budget more efficiently. For example, different allocation schemes of voting machines can be tested via simulation and the outcomes of these simulations used to decide which allocation scheme has the most desirable waiting time distribution for voters [Allen and Bernshteyn 2006]. The critical step in applying an R&S procedure to such a discrete simulation optimization application is to decide how many samples to draw from each candidate system. If we make this decision intelligently, we can find a good option more quickly and more reliably.

---

Part of this work is based on ideas introduced in the *Proceedings of the Winter Simulation Conference* under the title “Performance Measures for Ranking and Selection Procedures” [Waeber et al. 2010].

This work is supported in part by National Science Foundation grant CMMI-0800688 and the Air Force Office of Scientific Research FA9550-11-1-0083.

Authors’ addresses: R. Waeber (corresponding author), P. I. Frazier, and S. G. Henderson, 206 Rhodes Hall, Cornell University, Ithaca, NY 14853; email: [rw339@cornell.edu](mailto:rw339@cornell.edu).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permission may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2012 ACM 1049-3301/2012/08-ART16 \$15.00

DOI 10.1145/2331140.2331144 <http://doi.acm.org/10.1145/2331140.2331144>

A large number of R&S procedures have been introduced in the literature. Initially R&S problems were heavily studied in sequential statistical testing and experiment design (see Santner and Tamhane [1984] and Bechhofer et al. [1995] for overviews). In the recent simulation literature, selection procedures have become an integral component of simulation optimization algorithms [Boesel et al. 2003; Swisher et al. 2003]. Two theoretical frameworks have been introduced for evaluating the performance of R&S procedures. The first is the Indifference Zone (IZ) framework, initially introduced in [Bechhofer 1954], which considers the worst-case performance of an R&S policy over a set of configurations (see, for instance, Sobel and Huyett [1957], Paulson [1964], Rinott [1978], Kim and Nelson [2001], Kim and Nelson [2006a, 2006b]). The second is the Bayesian framework, which considers the average-case performance under a prior probability distribution over a set of configurations (see, for instance, Chick [1997] and Chick and Inoue [2001a, 2001b]). Each framework has advantages and disadvantages. The IZ framework provides a statistical guarantee on performance that provides peace of mind, but is often extremely conservative in the number of samples it requires an R&S procedure to take. The Bayesian framework, in contrast, results in R&S procedures that perform well in many cases, but might do badly in certain configurations. The two frameworks seem to be at odds with each other, presenting radically different approaches to the R&S problem at hand. While most existing R&S policies fall under one or the other of the two research directions, a class referred to as the Optimal Computing Budget Allocation (OCBA) schemes (see Chen and Lee [2010] and references within) encompasses aspects of both frameworks.

The nature of the aforesaid two research streams and the extensive set of available R&S policies raise a difficult question for practitioners: which R&S procedure should one use in an application? Branke et al. [2007] give empirically derived insights on how to select a selection procedure for independent normal R&S problems. They test existing policies on an extensive testbed and compare different performance measures. The goal of the current article is to provide a systematic machinery that one can use to select a selection procedure. We introduce a general theoretical framework that encompasses both the IZ and the Bayesian formulations of the R&S problem, as well as a continuum of formulations that strike a balance between the conservative worst-case analysis of IZ and the Bayesian emphasis on average-case performance. This framework helps us to choose among the large number of available R&S procedures in a systematic way. Surprisingly, the framework possesses a similar structure to that of convex risk measures studied in mathematical finance. See Föllmer and Schied [2004] and McNeil et al. [2005] for an introduction to risk measures used in finance.

We begin by describing how the performance of selection policies can be studied using a three-layer performance evaluation process. This extends statistical decision theory (see DeGroot [1970] and Berger [1985]) to the R&S setting. Choosing the three layers carefully allows the decision maker to determine what performance guarantees an R&S policy should satisfy for a given application. We define the *acceptance set* as the set of policies that meet these requirements. By allowing for flexible shapes of acceptance sets, the decision maker can express her risk-attitude towards different input scenarios (an input scenario assigns a risk weighting to each of the possible experiment configurations), and a trade-off between worst-case and average-case performance can be formulated in a geometric way. We state axioms an acceptance set should satisfy so that the induced preference order is consistent with intuition. We then introduce a new performance measure for R&S procedures solely based on the number of simulation iterations required for a policy to become acceptable. When performance can be controlled by the total number of simulation iterations applied, this performance measure is defined as the smallest simulation budget required for a policy to lie within the acceptance set. Total (expected) computational cost has been used to compare

performances of different R&S policies (see, for instance, Hong and Nelson [2005]). However, extensions to multiple input scenarios and higher-dimensional acceptance sets that allow for a robust performance assessment are, to the best of our knowledge, new in the R&S setting.

Through a concrete application, namely a Bernoulli R&S problem with an exit option appearing in the context of allocating voting machines, we investigate the behavior of different performance measures and their induced preference orders on a set of R&S procedures. Robust preference regions are identified and the suggested framework is used to determine optimal tuning parameters for R&S policies. This numerical example shows in detail how a decision maker can choose an R&S procedure that meets her individual performance requirements.

This article is based on ideas introduced in Waeber et al. [2010]. We provide a rigorous follow-up discussion on the properties of performance measures and further investigate the notion of acceptance sets. In addition, we guide the reader through a real-world example, which demonstrates how the suggested framework can be used in choosing among a set of R&S policies from a practitioner's perspective. More specifically, the main contributions of this article are as follows.

- We formulate an explicit three-layer performance analysis process for R&S procedures. This process extends ideas from statistical decision theory to R&S procedures and allows one to compare and optimize different R&S procedures.
- We investigate the connection to convex risk measures and the subsequent implications for performance measures of R&S procedures. In particular, the risk associated with the unknown underlying configuration of an experiment can be assessed in the same way as the risk associated with the unknown payoff of a financial portfolio.
- We introduce the idea of acceptance sets for R&S policies. This provides the decision maker with a tool to determine the performance guarantee a procedure should satisfy given her risk tolerance. Further, an efficient frontier on the set of R&S policies can be determined by solving an optimization problem over the set of acceptable policies.
- We demonstrate the practical benefits of our framework via a concrete application, namely the allocation of voting machines to precincts.

The outline of the article is as follows. Section 2 introduces the R&S problem and outlines the three-layer performance evaluation process. Section 3 analyzes the connection between existing performance measures and convex risk measures. Section 4 introduces acceptance sets and a new performance measure based on computational cost. Section 5 demonstrates how to use the suggested framework in a specific simulation-optimization application. Section 6 concludes.

## 2. PERFORMANCE EVALUATION FOR RANKING AND SELECTION PROCEDURES

In this section, we introduce the R&S problem, outline the performance evaluation process for R&S policies, and show how the IZ and the Bayesian formulation fit into this framework.

### 2.1. R&S Setting

We consider  $k$  systems. Each system  $i$  corresponds to a simulation experiment producing independent and identically distributed (iid) random variables  $Y_i \sim F_i \in \mathcal{F}_1$ ,  $i = 1, \dots, k$ , where  $F_i$  is the distribution of  $Y_i$  and  $\mathcal{F}_1$  is a family of 1-dimensional distributions. The random vector  $\mathbf{Y} := (Y_1, \dots, Y_k)$  has joint distribution  $\mathbf{F} \in \mathcal{F}$ , where  $\mathcal{F}$  is a family of  $k$ -dimensional distributions with marginal distributions belonging to

the family  $\mathcal{F}_1$ . When the  $k$  systems are simulated independently, the joint distribution is given as the product of the marginals, that is,  $\mathbf{F}(x_1, \dots, x_k) = F_1(x_1) \cdots F_k(x_k)$  for all  $(x_1, \dots, x_k) \in \mathbb{R}^k$ . Sometimes, for example, when common random numbers are used [Chick and Inoue 2001a; Nelson and Matejcik 1995], the outputs of the systems are not independent and  $\mathbf{F}$  incorporates a dependence structure. For this article we assume that the  $k$  systems are independent.

The performance of a single system is measured by a functional  $\theta$  on  $\mathcal{F}_1$ , that is,  $\theta_i := \theta(F_i) \in \mathbb{R}$ , and  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_k) \in \mathbb{R}^k$  summarizes the performance of all  $k$  systems. We are interested in identifying the system with the lowest performance parameter  $\theta_i$ , that is, we want to select system  $i^*$  where  $\theta_{i^*} = \min_{i \in \{1, \dots, k\}} \theta_i$ . The goal of identifying the system with the highest  $\theta_i$  can be accommodated by flipping the sign of  $\theta_i$ . Although  $\theta_i$  is often taken to be the mean of the distribution in the R&S literature, our framework allows other choices as well. For example, in finance, one is often interested in a portfolio composition with the highest Sharpe ratio,  $\theta_i = \mu_i / \sigma_i$ , where  $\mu_i$  is the expectation and  $\sigma_i$  the standard deviation of the daily returns  $Y_i$ . Because  $\boldsymbol{\theta}$  depends on  $\mathbf{F}$ , we sometimes write  $\boldsymbol{\theta}(\mathbf{F})$  to emphasize this dependence. More often, however, we simply write  $\boldsymbol{\theta}$  when the context is clear.

In general an R&S policy  $\pi = (x_{1:\tau}, \tau, i_\pi)$  consists of three components: an *allocation* rule  $x_{1:\tau}$  that determines which systems to simulate, a *stopping* rule  $\tau$  that determines when to stop simulating, and a *selection* rule  $i_\pi$  that determines which system to select based on the simulation results. We use the notation  $x_{1:n} := (x_1, \dots, x_n)$  for  $n \in \mathbb{N}$ . Formally, for  $n = 1, \dots, N$ , we define  $x_n \in \{1, \dots, k\}$  to be the system simulated at time  $n$  and  $Y_{x_n, n} \in \mathbb{R}$  the corresponding outcome. To simplify we often just write  $Y_n$  in place of  $Y_{x_n, n}$ . Conditioned on the choice  $x_n$ ,  $Y_n$  is distributed according to  $F_{x_n}$ , and is assumed independent of  $x_{1:n-1}$  and  $Y_{1:n-1}$ . Adaptive allocation rules are allowed, which means that the decision  $x_n$  can depend upon the previously observed outcomes, that is,  $x_n$  can be a function of  $x_{1:n-1}$  and  $Y_{1:n-1}$ , although we suppress this in the notation. In this article we restrict ourselves to deterministic stopping rules, that is,  $\tau \equiv N$  for  $N \in \mathbb{N}$ , which denotes the *total* number of measurements budgeted across all systems and  $N$  is specified at the beginning of the experiment. The restriction to deterministic stopping rules is not a requirement for the framework that we will introduce, but helps to simplify notation and discussion throughout the work (see Remark 1 for how the framework can be extended to policies where  $\tau$  is random). Define  $i_\pi \in \{1, \dots, k\}$  to be the selection made once all  $N$  simulations have been carried out. This decision is a function of all the observed outcomes, that is,  $i_\pi$  is a function of  $x_{1:N}$  and  $Y_{1:N}$ . Finally, let  $\Pi$  denote the set of all possible policies  $\pi$ .

The essential question in R&S is this: which policy  $\pi \in \Pi$  should we use?

## 2.2. Three-Layer Performance Evaluation

When we use an R&S procedure, we hope that it selects the best alternative. Inevitably, however, it will sometimes fail to do so. In this section, we discuss ways to measure the risk associated with this failure to select the best. Together, there are three concepts that determine this risk: the loss of the decision, the configuration-specific risk, and the overall risk.

**2.2.1. Loss of the Decision (Outcome-Specific Risk).** If  $\boldsymbol{\theta}$  is known, we can assign a loss quantity  $L(i_\pi, \boldsymbol{\theta}) \in \mathbb{R}$  to the decision  $i_\pi$  that reflects the *loss* (cost) associated with choosing system  $i_\pi$ . Examples of such loss functions  $L$  include the following.

- (1)  $L(i_\pi, \boldsymbol{\theta}) := \mathbb{1}\{\theta_{i_\pi} \neq \theta_{i^*}\}$  where  $\mathbb{1}\{\cdot\}$  denotes the indicator function that is 1 if its argument is true and 0 otherwise. This loss function is the 0-1 loss, used in much of

the R&S literature. It penalizes all incorrect selections equally. The corresponding expected loss is the *Probability of InCorrect Selection* (PICS).

- (2)  $L(i_\pi, \theta) := \mathbb{1}\{\theta_{i_\pi} \geq (\theta_{i^*} + \delta)\}$ , for  $\delta > 0$ . This loss function only penalizes incorrect selection when the selected alternative is more than  $\delta$  worse than the best. This reflects the idea that we may be indifferent to differences in performance of  $\delta$  or less. Such a notion is similar to the approach in Nelson and Banerjee [2001].
- (3)  $L(i_\pi, \theta) := \theta_{i_\pi} - \theta_{i^*}$ . This is the linear loss function used in [Chick and Inoue 2001b], and is also referred to as *regret*. It is equal to the difference in performance between the best system and the selected system.
- (4)  $L(i_\pi, \theta) := f(\theta_{i_\pi} - \theta_{i^*})$ , where  $f$  is a convex and increasing function on  $\mathbb{R}$ . This generalizes linear loss, with the function  $f$  modeling the risk aversion of the decision maker. Possible choices for  $f$  include  $f(x) = \mathbb{1}\{x \geq 0\}x^p$  for some constant  $p \geq 1$  and  $f(x) = \alpha \exp(\gamma x)$  for some constants  $\alpha, \gamma > 0$ .
- (5)  $L(i_\pi, \theta) := \theta_{i_\pi} - c$ , where  $c \in \mathbb{R}$  is a constant. This can be used when our loss depends not on the difference in performance between the selected and the best, but between the selected and some known threshold. In some cases, for example, the Bayesian framework using expected loss, this induces the same preference order as the linear loss function. In other cases, it does not.

**2.2.2. Configuration-Specific Risk.** The loss  $L(i_\pi, \theta)$  quantifies the risk associated with a single outcome of the selection decision  $i_\pi = i_\pi(x_{1:N}, Y_{1:N})$ . However, this decision  $i_\pi$  made by an R&S policy  $\pi$  depends on the random simulation outcomes, and so  $L(i_\pi, \theta)$  is itself a random variable even when  $\theta$  is known. Thus, an additional source of risk is the randomness induced by stochastic simulation, given a single fixed configuration  $\mathbf{F}$ . We call this *configuration-specific risk*, and refer to it as  $R(\pi, \mathbf{F})$ .

Every policy  $\pi$  and configuration  $\mathbf{F}$  induces a probability measure  $\mathbb{P}_{\mathbf{F}}^\pi$  on the sampling space. To quantify the configuration-specific risk, a functional of the random variable  $L(i_\pi, \theta)$  with respect to  $\mathbb{P}_{\mathbf{F}}^\pi$  needs to be determined. The most common choice is to use the expected loss, that is,  $R(\pi, \mathbf{F}) := \mathbb{E}_{\mathbf{F}}^\pi[L(i_\pi, \theta)]$ . This is a reasonable measure when the policy is used repeatedly for R&S in a large number of similar situations, or by a decision maker with little aversion to risk. However, when selection errors are costly, or if a policy is to be used only once or a small number of times (for example, in medical trials), more conservative functionals of the loss distribution, such as quantiles or expected loss above a given quantile, might be considered [Chick 1997]. Generally, we write the configuration-specific risk as

$$R(\pi, \mathbf{F}) := r(L(i_\pi, \theta)), \quad (1)$$

where  $r$  is a functional of the distribution of the random variable  $L(i_\pi, \theta)$  with respect to  $\mathbb{P}_{\mathbf{F}}^\pi$ , given a fixed  $\mathbf{F}$ . Defining both  $L$  and  $R$  allows some flexibility, for example, using the loss function  $L_1(i_\pi, \theta) = f(\theta_{i_\pi} - \theta_{i^*})$  for some convex and increasing function  $f$  introduced in the previous section and the configuration-specific risk  $R(\pi, \mathbf{F}) = \mathbb{E}_{\mathbf{F}}^\pi[L_1(i_\pi, \theta)]$  is equivalent to using the linear regret loss  $L_2(i_\pi, \theta) = \theta_{i_\pi} - \theta_{i^*}$  and the configuration-specific risk  $R(\pi, \mathbf{F}) = \mathbb{E}_{\mathbf{F}}^\pi[f(L_2(i_\pi, \theta))]$ .

**2.2.3. Overall Risk.** The configuration-specific risk  $R(\pi, \mathbf{F})$  defined in (1) still depends on the *unknown* underlying configuration  $\mathbf{F} \in \mathcal{F}$ . There is also risk associated with the fact that this configuration is unknown. Especially when an R&S procedure is used repeatedly over varying configurations  $\mathbf{F}$  the risk of unknown  $\mathbf{F}$  needs to be treated separately from the functional  $R(\pi, \mathbf{F})$ . A risk-averse decision maker would prefer a policy  $\pi$  that has moderately low  $R(\pi, \mathbf{F})$  across all possible configurations  $\mathbf{F}$ . In contrast, a decision maker indifferent to extreme risks might prefer a policy  $\pi$  with low  $R(\pi, \mathbf{F})$  in most configurations, but high  $R(\pi, \mathbf{F})$  in a few problematic configurations.



To quantify the overall risk, we use a functional  $\rho$  on the mapping  $\mathbf{F} \mapsto R(\pi, \mathbf{F})$ . With a slight abuse of notation, we write  $\rho(\pi) = \rho(R(\pi, \cdot)) \in \mathbb{R}$ . This performance mapping  $\rho$  depends only on the policy  $\pi$ , and not on the actual configuration  $\mathbf{F}$ . It can then be used to compare different policies and, together with the natural ordering of  $\mathbb{R}$ , induces the sought-after *preference order* on  $\Pi$ . A decision maker *prefers* policy  $\pi^{(1)}$  over policy  $\pi^{(2)}$ , that is,  $\pi^{(1)} \succ \pi^{(2)}$ , if and only if  $\rho(\pi^{(1)}) < \rho(\pi^{(2)})$ . A decision maker is *indifferent* between policy  $\pi^{(1)}$  and  $\pi^{(2)}$ , that is,  $\pi^{(1)} \sim \pi^{(2)}$  if and only if  $\rho(\pi^{(1)}) = \rho(\pi^{(2)})$ .

Depending on the risk tolerance of the decision maker there are different ways to define  $\rho$ . The three most popular choices are.

- (1) *Worst-Case (WC) performance*:

$$\rho_{WC}(\pi) := \sup_{\mathbf{F} \in \mathcal{F}} R(\pi, \mathbf{F}).$$

This performance measure prefers policies which perform well in the worst underlying configuration of the experiment. This is the most conservative performance measure.

- (2) *Indifference Zone (IZ)*:

$$\rho_{IZ}(\pi) := \sup_{\mathbf{F} \in \mathcal{F} \setminus IZ} R(\pi, \mathbf{F}). \quad (2)$$

The IZ approach was first introduced in the seminal work [Bechhofer 1954], and since then the IZ approach has become a standard formulation in the R&S literature. This formulation is closely related to the worst case performance measure, being the worst-case over all but a subset (the *indifference-zone*)  $IZ \subset \mathcal{F}$  of configurations. The IZ performance measure prefers policies that perform well in the worst underlying configuration *outside* the indifference zone. Inside the indifference zone, configurations are assumed to have alternatives with performances  $\theta_i$  so similar that it is not worth detecting the difference. This formulation is still quite conservative, but is not as conservative as WC.

- (3) *Bayes risk*. The Bayesian approach assumes the existence of a  $\sigma$ -field  $\mathcal{H}$  on  $\mathcal{F}$  and incorporates a prior probability measure  $\mathbb{P}_0$  on the measurable space  $(\mathcal{F}, \mathcal{H})$ . This prior reflects a *risk weighting* on  $\mathbf{F}$  and need not correspond to the prior belief regarding  $\mathbf{F}$  or  $\theta$ . In the Bayesian approach,  $\mathbf{F}$  is considered random and  $R(\pi, \cdot)$  is assumed measurable, that is, a random variable. Different functionals of the distribution of  $R(\pi, \mathbf{F})$  with respect to  $\mathbb{P}_0$  can be used to quantify the performance of a policy. Again, the most popular choice is to use the expectation, that is,

$$\rho_{Bayes}(\pi) := \mathbb{E}_{\mathbb{P}_0}[R(\pi, \mathbf{F})] = \int_{\mathcal{F}} R(\pi, \mathbf{F}) \mathbb{P}_0(d\mathbf{F}).$$

In contrast to the worst-case and the indifference zone formulations, the Bayes formulation focuses on an average-case performance rather than a conservative worst-case performance. When the risk weighting  $\mathbb{P}_0$  is diffuse, this average-case analysis might be too optimistic. One can introduce more risk aversion by using a more conservative risk weighting  $\mathbb{P}_0$  or a different functional of  $R(\pi, \mathbf{F})$  with respect to  $\mathbb{P}_0$ , for example, quantiles or expected shortfall.

**2.2.4 Decision-Theoretic Formulation for R&S Procedures.** To summarize, in comparing different policies, three quantities must be defined:

- (1) the loss of a decision  $L(i_\pi, \theta)$ ;
- (2) the configuration-specific risk  $R(\pi, \mathbf{F})$ ;
- (3) the overall risk  $\rho(\pi)$ .

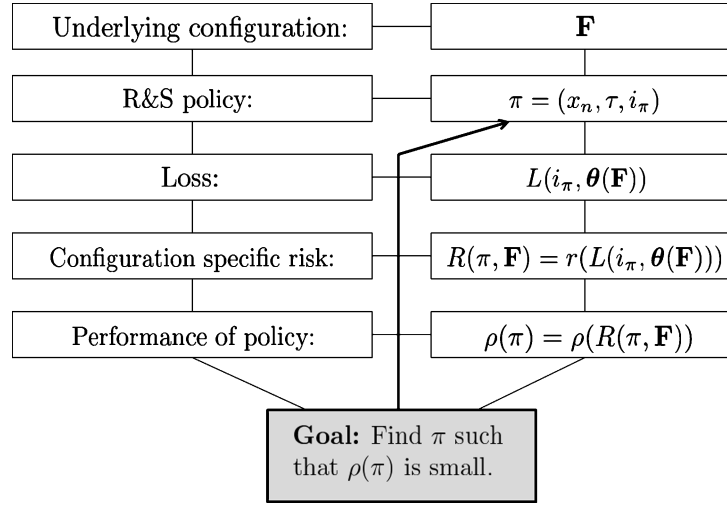


Fig. 1. Necessary quantities to define a preference order on the set of R&S policies.

With these quantities defined, policy  $\pi^{(1)}$  is preferred over policy  $\pi^{(2)}$  if and only if  $\rho(\pi^{(1)}) < \rho(\pi^{(2)})$ . Figure 1 provides an overview of these three concepts and how they interact. There is no single best way to define these quantities. Rather, they should be chosen carefully according to the decision maker's risk preferences for the particular problem instance addressed. In Section 5 we explain how these three levels can be determined for a concrete application.

### 3. ROBUST REPRESENTATION OF PERFORMANCE MEASURES

In the previous section, we described the process that underlies the performance evaluation of R&S procedures. We now focus on a unifying representation of the three popular performance measures defined in Section 2.2.3. It turns out that this representation is analogous to the robust representation of *convex risk measures* studied in mathematical finance. This robust representation then suggests other performance measures between the popular measures given in Section 2.2.3 that may overcome their drawbacks.

In financial applications a risk measure should properly quantify the risk of a financial portfolio. For this, the return of a portfolio over a given time horizon, usually 10 days, is modeled as a random variable  $X$ . Let  $L^0$  denote the space of random variables which are almost surely finite. In finance, a *risk measure* is a mapping  $\chi : L^0 \rightarrow \mathbb{R}$ . The seminal work Markowitz [1952] on portfolio selection used the variance of  $X$  as a risk measure. The variance, however, has certain drawbacks: it lacks monotonicity and cash invariance (see what follows). Other popular risk measures, such as Value-at-Risk, may penalize diversification of a portfolio. These drawbacks initiated a discussion of the structural properties that a risk measure should satisfy in financial applications [Artzner et al. 1999]. The class of *convex risk measures*, introduced in Föllmer and Schied [2002], has emerged as a class of good risk measures for financial applications.

**Definition 3.1.** A mapping  $\chi : L^0 \rightarrow \mathbb{R}$  is called a *convex risk measure* if it satisfies the following conditions for all  $X, Y \in L^0$ .

- (1) *Monotonicity.* If  $X \leq Y$ , then  $\chi(X) \geq \chi(Y)$ . That is, a portfolio that always has a smaller payoff than the other portfolio should be considered riskier.

- (2) *Cash Invariance.* If  $m \in \mathbb{R}$ , then  $\chi(X + m) = \chi(X) - m$ . That is, if  $m$  dollars are added to a portfolio, its risk should decrease by  $m$  dollars. This property is important in determining appropriate risk capital for financial institutions. For example, regulatory agencies may require that the cash position  $m$  is large enough so that the risk of a portfolio is equal to 0.
- (3) *Convexity.*  $\chi(\lambda X + (1 - \lambda)Y) \leq \lambda\chi(X) + (1 - \lambda)\chi(Y)$ , for  $0 \leq \lambda \leq 1$ . That is, a diversified portfolio should not be considered riskier than an undiversified portfolio.

At a first glance, convex risk measures do not seem related to R&S. However, a strong connection can be identified through the following robust representation property of convex risk measures.

**THEOREM 3.2** [FÖLLMER AND SCHIED 2002]. *Denote by  $\mathcal{Q}$  the set of all finitely additive probability measures on a measurable space  $(\Omega, \mathcal{S})$ . Let  $\alpha : \mathcal{Q} \rightarrow \mathbb{R} \cup \{\infty\}$  be a functional with  $\inf_{Q \in \mathcal{Q}} \alpha(Q) \in \mathbb{R}$  and  $X \in L^\infty$  ( $L^\infty$  denotes the space of random variables that are almost surely bounded). Then the mapping*

$$\chi(X) := \sup_{Q \in \mathcal{Q}} (\mathbb{E}_Q[-X] - \alpha(Q)), \quad (3)$$

*defines a convex risk measure on  $L^\infty$ .*

With further technical assumptions the reverse direction also holds, that is, every convex risk measure has a representation of the form (3). See Föllmer and Schied [2002, 2004] for more details on convex risk measures used in finance and a proof of Theorem 3.2.

We now state a similar representation result for the performance measures used for R&S procedures. (We denote with  $\mathcal{B}(\mathbb{R})$  the Borel  $\sigma$ -field on  $\mathbb{R}$ .)

**PROPOSITION 3.3.** *Consider a loss function  $L$  and a configuration-specific risk mapping  $R$ . If the space  $\mathcal{F}$  has a  $\sigma$ -field  $\mathcal{H}$  such that for a given policy  $\pi$  the mapping  $R(\pi, \cdot) : (\mathcal{F}, \mathcal{H}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$  is measurable, then the worst-case, indifference zone, and Bayesian risk for Ranking and Selection procedures defined in Section 2.2.3 can each be represented as*

$$\rho(\pi) = \sup_{Q \in \mathcal{Q}} (\mathbb{E}_Q[R(\pi, \mathbf{F})] - \alpha(Q)), \quad (4)$$

*where  $\mathcal{Q}$  is some appropriate set of probability measures on  $(\mathcal{F}, \mathcal{H})$ , and  $\alpha : \mathcal{Q} \rightarrow \mathbb{R} \cup \{\infty\}$  is some penalty function.*

**PROOF.** Let  $\alpha \equiv 0$ . Then

- the set  $\mathcal{Q}_{WC} = \{\text{all Dirac point measures for } \mathbf{F} \in \mathcal{F}\}$  yields  $\rho(\pi) = \rho_{WC}(\pi)$ ;
- the set  $\mathcal{Q}_{IZ} = \{\text{all Dirac point measures for } \mathbf{F} \in \mathcal{F} \setminus IZ\}$  yields  $\rho(\pi) = \rho_{IZ}(\pi)$ ;
- the set  $\mathcal{Q}_{Bayes} = \{P_0\}$  yields  $\rho(\pi) = \rho_{Bayes}(\pi)$ .

A Dirac point measure for  $\mathbf{F} \in \mathcal{F}$  is a mapping  $\delta_{\mathbf{F}} : \mathcal{H} \rightarrow \{0, 1\}$ , such that  $\delta_{\mathbf{F}}(H) = 1$  if  $\mathbf{F} \in H$  and  $\delta_{\mathbf{F}}(H) = 0$  if  $\mathbf{F} \notin H$ .  $\square$

In this proposition,  $Q$  is the probability measure on the measurable space  $(\mathcal{F}, \mathcal{H})$ . This specifies the idea that the underlying experiment configuration  $\mathbf{F}$  is considered as a random realization from  $\mathcal{F}$ . The measurability of  $R(\pi, \cdot)$  is necessary so that  $\mathbb{E}_Q[R(\pi, \mathbf{F})]$  is defined. If, for example,  $\mathcal{F}$  belongs to a parametric family then this assumption is always satisfied by choosing  $\mathcal{H}$  as the corresponding Borel  $\sigma$ -field on the parameter set.



There is a sign change inside the expectations between representation (3) and (4). This is because, in finance  $X$  represents a return and hence  $-X$  represents the loss, whereas in the R&S setting the quantity  $R(\pi, \mathbf{F})$  refers directly to the “loss” of policy  $\pi$  under realization  $\mathbf{F}$ .

Proposition 3.3 raises two questions about performance measures for R&S policies that are worth investigating. First, are the properties of convex risk measures necessary for the evaluation of R&S policies? Second, does the class of convex risk measures provide other useful performance measures beyond WC, IZ, and Bayes performance?

### 3.1. Properties of Risk Measures for R&S Procedures

In mathematical finance the properties of convex risk measures, that is, monotonicity, cash invariance, and convexity, are motivated by economic principles of investing. The random variable  $X$  is the unknown payoff of a financial portfolio over a given time period and  $\rho(X)$  the associated risk of such a portfolio. Proposition 3.3 suggests that the risk of an R&S procedure can be assessed in a similar way, where  $R(\pi, \mathbf{F})$  is the unknown and undesirable “payoff” and  $\rho(\pi)$  quantifies the risk of  $\pi$ . It is instructive to investigate what the properties of convex risk measures imply in the R&S setting.

**3.1.1. Monotonicity.** Monotonicity ensures that, if for every configuration  $\mathbf{F}$  policy  $\pi^{(1)}$  has higher configuration-specific risk than policy  $\pi^{(2)}$ , then the overall risk of  $\pi^{(1)}$  is also higher than the overall risk of policy  $\pi^{(2)}$ . This means that the induced preference order should never favor a policy that is *dominated* by another policy, where dominated is meant in the classical sense from statistical decision theory (see [Berger 1985]). Monotonicity is a property that *every* good performance measure for R&S procedures should satisfy.

An example of a nonmonotone performance measure for R&S is the *variance* of  $R(\pi, \cdot)$  for risk weighting  $Q$ , that is,  $\rho(\pi) := \text{Var}_Q(R(\pi, \mathbf{F}))$ . Comparing the performance of R&S procedures using the variance could cause noncoherent preference orders. Consider, for example, a Bernoulli R&S problem with  $k = 2$  alternatives and the goal of finding the system with the smallest success probability. Using the 0-1 loss function  $L(i_\pi, \theta) = \mathbb{1}\{\theta_i \neq \theta_{i^*}\}$  and the expected loss as configuration-specific risk,  $R(\pi, \mathbf{F})$  is the Probability of InCorrect Selection (PICS). For Bernoulli R&S the underlying configuration  $\mathbf{F}$  can be parameterized by the vector of success probabilities,  $\theta = (\theta_1, \theta_2)$ , where  $\theta_i$  is the success probability of system  $i$  for  $i = 1, 2$ . We compare two policies: policy  $\pi^{(1)}$  selects with equal probability system 1 or system 2; and policy  $\pi^{(2)}$  simulates each system independently once and selects the system with a failure ( $Y_i = 0$ ) as the better system. In case of a tie ( $Y_1 = Y_2$ ),  $\pi^{(2)}$  selects system 1 or 2 with equal probability. The configuration-specific risk for the two policies are:

$$R(\pi^{(1)}, \mathbf{F}) = \begin{cases} 0.5, & \text{if } \theta_1 \neq \theta_2, \\ 0, & \text{if } \theta_1 = \theta_2, \end{cases}$$

$$R(\pi^{(2)}, \mathbf{F}) = \begin{cases} 0.5(1 + \theta_1 - \theta_2), & \text{if } \theta_1 < \theta_2, \\ 0, & \text{if } \theta_1 = \theta_2, \\ 0.5(1 + \theta_2 - \theta_1), & \text{if } \theta_1 > \theta_2. \end{cases}$$

This shows that  $R(\pi^{(1)}, \mathbf{F}) \geq R(\pi^{(2)}, \mathbf{F})$  for all  $\mathbf{F} \in \mathcal{F}$  and  $R(\pi^{(1)}, \mathbf{F}) > R(\pi^{(2)}, \mathbf{F})$  for  $\mathbf{F}$  with  $\theta_1 \neq \theta_2$ . If we assume a risk weighting  $Q$  with positive density on  $[0, 1]^2$  then  $\text{Var}_Q(R(\pi^{(1)}, \mathbf{F})) = 0 < \text{Var}_Q(R(\pi^{(2)}, \mathbf{F}))$ , so the variance as a performance measure would prefer policy  $\pi^{(1)}$ , even though  $\pi^{(1)}$  is dominated by  $\pi^{(2)}$ .

**3.1.2. Cash Invariance.** In mathematical finance cash invariance ensures that the risk measure  $\rho(X)$  can be used to determine appropriate capital requirements. If a portfolio manager has set aside a risk capital of the size  $\rho(X)$  then his portfolio is considered “safe” and hence *acceptable*.

For R&S procedures a similar interpretation is useful when the goal of using an R&S procedure is to minimize opportunity cost, as for example in Chick and Inoue [2001b], Chick and Gans [2009], Chick and Frazier [2009]. For a decision maker using policy  $\pi$ , the configuration-specific risk  $R(\pi, \mathbf{F})$  is her monetary penalty when the corresponding configuration is  $\mathbf{F}$ . The performance measure  $\rho(\pi)$  can then be used to calculate the capital this decision maker needs to set aside for future payments. If  $\rho(\pi)$  is negative, the decision maker can safely withdraw cash. With this interpretation cash invariance is a reasonable requirement on a performance measure for some R&S settings.

Acceptable portfolios and the determination of sufficient risk capital in a financial setting have another interesting parallel in the R&S setting. The invested computational effort for an R&S procedure can be interpreted as the allocated risk capital of a decision maker. The minimal required simulation budget for a policy  $\pi$  to be considered “safe” or acceptable provides the decision maker with a tool to assign the proper amount of computational effort to the R&S procedure. This approach is further investigated in Section 4 and leads to the notion of acceptable policies for R&S procedures.

**3.1.3. Convexity.** Convexity is an important property in the financial setting because it states that a diversified portfolio should not be riskier than an undiversified portfolio. The interpretation of the convexity property for R&S policies, that is, for  $0 \leq \lambda \leq 1$ ,

$$\rho(\lambda R(\pi^{(1)}, \cdot) + (1 - \lambda)R(\pi^{(2)}, \cdot)) \leq \lambda \rho(R(\pi^{(1)}, \cdot)) + (1 - \lambda)\rho(R(\pi^{(2)}, \cdot)), \quad (5)$$

is less obvious. The left side of (5) does *not* correspond to the risk of a policy that assigns  $\lambda N$  of the simulation budget to policy  $\pi^{(1)}$  and  $(1 - \lambda)N$  of the simulation budget to policy  $\pi^{(2)}$  nor to the risk of a randomized policy that applies policy  $\pi^{(1)}$  with probability  $\lambda$  and policy  $\pi^{(2)}$  with probability  $1 - \lambda$ .

We provide a possible interpretation of convexity for the R&S setting in the context of parallel computing. Consider the following simplified example: Two processors must perform R&S on the same sequence of configurations but cannot communicate with each other. Two R&S policies  $(\pi^{(1)}, \pi^{(2)})$  are available. We assume that the space  $\mathcal{F}$  can be partitioned into two nonempty sets  $A_1$  and  $A_2$  such that  $R(\pi^{(i)}, \mathbf{F}) = (1 - \mathbb{1}\{\mathbf{F} \in A_i\})$  for  $i = 1, 2$ , that is,  $\pi^{(1)}$  performs better when  $\mathbf{F} \in A_1$  and  $\pi^{(2)}$  performs better when  $\mathbf{F} \in A_2$ . If there are two experiment configurations in the sequence, then a procedure  $\bar{\pi}$  that uses policy  $\pi^{(1)}$  on one processor and policy  $\pi^{(2)}$  on the other processor for each configuration has total risk  $\rho(0.5R(\pi^{(1)}, \cdot) + 0.5R(\pi^{(2)}, \cdot))$ , that is, the left side of (5), whereas a procedure  $\tilde{\pi}$  that uses policy  $\pi^{(1)}$  for the first configuration and policy  $\pi^{(2)}$  for the second configuration has total risk  $0.5\rho(R(\pi^{(1)}, \cdot)) + 0.5\rho(R(\pi^{(2)}, \cdot))$ , that is, the right side of (5). Here  $\lambda$  is set to 0.5 since we only have two processors. If we should have infinitely many processors, any  $\lambda \in [0, 1]$  can be used. The convexity property on  $\rho$  assures that the procedure  $\bar{\pi}$  is not preferred over the procedure  $\tilde{\pi}$ . This is a desirable property because  $\tilde{\pi}$  provides a more stable penalty payment than  $\bar{\pi}$  and is considered less risky. In other words, the diversified policy  $\tilde{\pi}$  assures risk reduction in the same sense as diversification reduces the risk of a financial portfolio.

The preceding discussion of convex risk measures for R&S procedures suggests that the class of convex risk measures, for example, mappings that satisfy the representation (4), should be used to compare R&S policies in certain settings. For other R&S settings the cash invariance and convexity property might be less relevant and one could use a performance measure outside the class of convex risk measures,

for example, a quantile of  $R(\pi, \mathbf{F})$ . However, the monotonicity property should always be satisfied by a performance measure that is used to compare R&S policies.

### 3.2. Alternative Performance Measures

Proposition 3.3 shows that the three most popular risk measures for R&S procedures are convex risk measures. This suggests that the class of convex risk measures, that is, the class of mappings with the representation (4), is helpful in analyzing the performance of R&S policies. The following examples show how the representation (4) can construct risk measures that provide a robust risk assessment of R&S policies and with this overcome some of the disadvantages of existing performance measures.

- A decision maker might put emphasis on average-case performance, but also be somewhat concerned about IZ performance. She might take  $\mathcal{Q} = \mathcal{Q}_{IZ} \cup \mathcal{Q}_{Bayes}$ ,  $\alpha(Q) = 0$  for  $Q \in \mathcal{Q}_{Bayes}$ , and  $\alpha(Q)$  equal to a large strictly positive constant  $c$  for  $Q \in \mathcal{Q}_{IZ}$ . With this large value of the penalty function on  $\mathcal{Q}_{IZ}$ , the performance measure  $\rho(\pi)$  would be equal to the Bayes performance unless the worst-case performance was extremely bad, when  $\rho(\pi)$  would equal the worst-case performance minus  $c$ .
- Another decision maker might be interested only in the performance of a policy under the Bayes formulation, but be unsure about robustness of the performance with respect to the risk weighting  $\mathbb{P}_0$ . She could then use  $\mathcal{Q}_C = \{\mathbb{P} : \mathbb{P} = (1 - \varepsilon)\mathbb{P}_0 + \varepsilon\mathbb{C}, \mathbb{C} \in \mathcal{C}\}$ , for  $0 < \varepsilon < 1$  and a class of possible *contaminations*  $\mathcal{C}$ . Such a class of contaminated priors is studied in Bayesian robustness theory [Berger 1985].

## 4. ACCEPTANCE SETS AND A NEW PERFORMANCE MEASURE

In the previous sections we described how, given a fixed simulation budget  $N$ , the performance of a policy  $\pi$  can be evaluated using a performance measure  $\rho$ .

This allows us to compare two policies using the same simulation budget. In practice, one can control the total simulation budget. In this section we describe a class of performance evaluations that take this into account. In this description we first specify what it means for a policy to be acceptable. We then quantify the performance of a policy as the smallest simulation budget required to make it acceptable. Acceptability is defined in terms of acceptance sets, and by allowing acceptance sets with different shapes, we allow for a range of trade-offs between worst-case and average-case performance.

To motivate the concept of acceptable policies, we use the representation (4) of  $\rho$  in Proposition 3.3. While the IZ and Bayesian formulations are special cases of (4), this representation can be used to define much more general performance measures as seen in Section 3.2. Each probability measure  $Q$  corresponds to a *scenario* of the underlying configuration  $\mathbf{F}$ , that is, a distribution over configurations, and  $\mathcal{Q}$  represents the set of all scenarios considered. One interpretation is that a scenario  $Q$  can be seen as one possible prior belief on the underlying configuration. Considering a set of scenarios  $\mathcal{Q}$  implies our uncertainty about this prior belief. Examples of  $\mathcal{Q}$  include  $\mathcal{Q}_{WC}$ ,  $\mathcal{Q}_{IZ}$ ,  $\mathcal{Q}_{Bayes}$ , and  $\mathcal{Q}_C$  as defined in Section 3. From now on we assume that  $|\mathcal{Q}| = m < \infty$ . The extension to infinitely many scenarios is mathematically more challenging and left open for future research.

Given a loss function  $L(i_{\pi_N}, \theta)$ , a risk functional  $R(\pi_N, \mathbf{F})$ , and a penalty function  $\alpha$ , the risk of a policy  $\pi_N$  that takes  $N$  measurements can be represented as a mapping  $f(\pi_N, \cdot) : \mathcal{Q} \rightarrow \mathbb{R}$  defined by  $f(\pi_N, Q) = \mathbb{E}_Q[R(\pi_N, \mathbf{F})] - \alpha(Q)$ . (To emphasize the dependence on  $N$  we now write  $\pi_N$  instead of just  $\pi$ .) The set of *all* risks generated by arbitrary combinations of  $L(\cdot, \theta)$ ,  $R(\cdot, \mathbf{F})$ ,  $\alpha$ , and  $N$  builds a subspace of the functional space  $\mathcal{G} = \{f : \mathcal{Q} \rightarrow \mathbb{R}\}$ . The space  $\mathcal{G}$  is equivalent to  $\mathbb{R}^m$  (recall that  $|\mathcal{Q}| = m < \infty$ , so

each function mapping  $\mathcal{Q}$  to  $\mathbb{R}$  can be written as a vector with  $m$  elements). Thus, the set of *attainable* risks corresponds to a subset of  $\mathbb{R}^m$ . The risk of a particular policy  $\pi_N$  can then be represented as a risk vector

$$f(\pi_N, \mathcal{Q}) := \begin{pmatrix} f(\pi_N, Q_1) \\ \vdots \\ f(\pi_N, Q_m) \end{pmatrix} = \begin{pmatrix} \mathbb{E}_{Q_1}[R(\pi_N, \mathbf{F})] - \alpha(Q_1) \\ \vdots \\ \mathbb{E}_{Q_m}[R(\pi_N, \mathbf{F})] - \alpha(Q_m) \end{pmatrix} \in \mathbb{R}^m.$$

We then suppose that the decision maker specifies an *acceptance* set  $\mathcal{A} \subseteq \mathbb{R}^m$  such that a policy  $\pi_N$  is *acceptable* if and only if  $f(\pi_N, \mathcal{Q}) \in \mathcal{A}$ . Such an acceptance set represents the decision maker's risk-attitude towards different scenarios  $Q_i$ ,  $i = 1, \dots, m$ , and should be defined in a coherent way, that is, any reasonable acceptance set should satisfy certain axioms. We assume that  $f$  is normalized so that,  $f(\pi_N, Q) \in [0, 1]$  for all  $Q$  in  $\mathcal{Q}$ , where  $f(\pi_N, Q) = 0$  corresponds to the minimal risk and  $f(\pi_N, Q) = 1$  to the maximal risk. These minimal and maximal risks do not need to be achieved by any given  $\pi_N$ . Further, we assume that the policy that chooses one of the  $k$  systems uniformly at random has risk  $(k-1)/k$ . The normalization corresponds specifically to the Probability of InCorrect Selection (PICS) formulation. If a different loss function  $L$  or a different configuration-specific risk  $R$  is used the normalization may change as well. We then require that any *coherent* acceptance set  $\mathcal{A}$  should satisfy the following axioms. (In the following we use vector notation to define rectangular subsets of  $\mathbb{R}^m$ , for example, for two vectors  $\mathbf{p}, \mathbf{q} \in \mathbb{R}^m$ , we define  $[\mathbf{p}, \mathbf{q}] := [p_1, q_1] \times [p_2, q_2] \times \dots \times [p_m, q_m] \subseteq \mathbb{R}^m$ .)

- Axiom 1.  $\mathbf{0} \in \mathcal{A}$ . A policy with no risk under any scenario should always be acceptable.
- Axiom 2.  $((k-1)/k \cdot \mathbf{1}, \mathbf{1}] \cap \mathcal{A} = \emptyset$ , where  $\mathbf{1}$  represents a vector of ones in  $\mathbb{R}^m$ . Any acceptable policy should be as good as randomly selecting a “best system.” Note that the formulation of this axiom depends on our choice of normalization.
- Axiom 3. For any point  $\mathbf{q} \in \mathcal{A}$ ,  $[\mathbf{0}, \mathbf{q}] \subseteq \mathcal{A}$ . This axiom ensures monotonicity on the set of policies, that is, if a policy dominates an acceptable policy then this policy should also be acceptable.

It follows from the third axiom that if  $\mathbf{q} \notin \mathcal{A}$ , then  $[\mathbf{q}, \mathbf{1}] \cap \mathcal{A} = \emptyset$ . Any policy dominated by an unacceptable policy should also be unacceptable.

Once a decision maker has identified an acceptance set according to her risk tolerance, we can check whether or not a policy  $\pi_N$  that takes  $N$  measurements is acceptable or not. Then, instead of predetermining the simulation budget  $N$ , we can treat  $N$  as a control for the policy  $\pi$ . This interpretation motivates a new performance measure

$$\varphi_{\mathcal{A}}(\pi) := \inf \{N \in \mathbb{N} \mid f(\pi_N, \mathcal{Q}) \in \mathcal{A}\}, \quad (6)$$

which is the minimum simulation budget required to make the policy  $\pi$  acceptable to the decision maker.

The intuition behind this performance measure is that the *cost* of performing R&S can be measured as the simulation time required to achieve a certain confidence level. If an infinite number of simulations is allowed, then any *consistent* policy would be satisfactory. Here a consistent policy is one that converges to the right answer, that is,  $f(\pi_N, \mathcal{Q}) \rightarrow \mathbf{0}$ , as  $N \rightarrow \infty$ . Therefore, a performance measure for varying  $N$  needs to depend on how quickly  $f(\pi_N, \mathcal{Q})$  converges to  $\mathbf{0}$ .

The conceptual plot in Figure 2 shows how the risk measure  $\varphi_{\mathcal{A}}$  can be visualized for two policies  $\pi^{(1)}$  and  $\pi^{(2)}$ , where the grey area is a possible acceptance set  $\mathcal{A}$ . In order to visualize the trajectories we consider only two scenarios,  $Q_1$  and  $Q_2$ . As  $N$  varies between 0 and  $\infty$ , each policy traces out a trajectory in  $\mathbb{R}^2$  indicated by the




Fig. 2. Possible acceptance set  $\mathcal{A}$  and trajectories of two policies  $\pi^{(1)}$  and  $\pi^{(2)}$ . The trajectories trace  $f(\pi_N, \mathcal{Q})$  as  $N \rightarrow \infty$ . The trajectories are not continuous curves, as  $N \in \mathbb{N}$ , but we draw them continuously here for simplicity.

dashed arrows. When a trajectory crosses the boundary of the acceptance set  $\mathcal{A}$  for the first time, the corresponding number of iterations  $N$  is precisely the performance of the corresponding policy.

The IZ and Bayesian formulations can be combined effectively in this framework. Consider scenarios  $Q_1$  as a point mass at the *Least Favorable Configuration* (LFC) (choose an arbitrary LFC if there exist a family of LFCs) and  $Q_2 = \mathbb{P}_0$ , some prior distribution on  $(\mathcal{F}, \mathcal{H})$ . The LFC is a configuration  $\mathbf{F}$  of the experiment where the supremum in (2) is (assumed to be) attained. Figure 3 shows the acceptance set  $\mathcal{A}_1$  of a decision maker using the IZ formulation, as well as the acceptance set  $\mathcal{A}_2$  of a decision maker using the Bayesian formulation. A decision maker who wants to apply both the IZ and the Bayesian formulation would then use the intersection of  $\mathcal{A}_1$  and  $\mathcal{A}_2$  as her acceptance set. The dashed arrows again indicate trajectories of two R&S policies.

It is instructive to analyze the connection between the fixed- $N$  performance measure  $\rho(\pi) = \sup_{Q \in \mathcal{Q}} (\mathbb{E}_Q[R(\pi, \theta)] - \alpha(Q))$  as defined in Proposition 3.3 and the variable- $N$  performance measure  $\varphi_{\mathcal{A}}(\pi) = \inf \{N \in \mathbb{N} | f(\pi_N, \mathcal{Q}) \in \mathcal{A}\}$  associated with an acceptance set  $\mathcal{A}$  as introduced in (6). The concept of acceptance sets in fact extends the representation (4), which implicitly, through the supremum operator, assumes rectangular acceptance sets  $[0 \cdot \mathbf{1}, d \cdot \mathbf{1}]$  for some constant  $d \in \mathbb{R}_+$ . That is, for a constant  $d \in \mathbb{R}$  defining the largest acceptable value of  $\rho(\pi)$ ,  $\rho(\pi_N) \leq d$  if and only if  $\mathbb{E}_Q[R(\pi_N, \mathbf{F})] - \alpha(Q) \leq d$  for each  $Q \in \mathcal{Q}$ , which holds if and only if  $f(\pi_N, \mathcal{Q})$  is in the set  $\mathcal{A} = [0 \cdot \mathbf{1}, d \cdot \mathbf{1}]$ .

While  $\rho(\pi)$  in (4) focuses on the worst-case performance over different scenarios  $Q \in \mathcal{Q}$ , choosing nonrectangular acceptance sets allows the decision maker to move away from a conservative worst-case analysis across scenarios, as in the numerical example in Section 5. The new performance measure  $\varphi_{\mathcal{A}}(\pi)$  then serves two purposes. First, it relaxes the conservative supremum operator of (4) but still induces a preference order on  $\Pi$  via the natural ordering of  $\mathbb{N}$ . Second, in many cases it is helpful to consider the performance of a policy under different acceptance sets. For example, when a decision maker uses an R&S procedure that is acceptable in terms of average-case performance, it is informative to determine the additional simulation budget required so that, in addition, an IZ performance guarantee is met. To this end, the performance measure  $\varphi_{\mathcal{A}}(\pi)$  as the number of “necessary” iterations gives an easy-to-interpret comparison between different acceptance sets for a policy  $\pi$ .






Fig. 3. Acceptance sets corresponding to the IZ formulation ( $\mathcal{A}_1$ ) and to the Bayesian formulation ( $\mathcal{A}_2$ ). The acceptance sets are chosen such that the performance requirement on the Bayesian scenario is stricter than the requirement on the IZ scenario. This reflects the typical case where the IZ (worst-case) performance is much worse than the Bayesian (average-case) performance. The intersection  $\mathcal{A}_1 \cap \mathcal{A}_2$  defines then a performance measure  $\varphi_{\mathcal{A}_1 \cap \mathcal{A}_2}$  that has a strong requirement on the Bayesian performance and guarantees that the IZ performance is not extremely bad.

*Remark 1.* So far we have assumed that the total number of simulation iterations  $N$  is deterministic. Simulation analysts, however, often use R&S procedures that have random run lengths. Indeed, it has been empirically shown [Branke et al. 2007] that, in many concrete examples, policies with adaptive simulation budgets outperform policies with fixed simulation budgets. Furthermore, the computational cost of an R&S procedure may depend on more than just the total number of iterations  $N$ , because simulation time may vary across alternative systems, and computations performed by the policy in choosing which systems to simulate may also contribute some cost. Hence the number of simulation iterations  $N$  may not appropriately reflect the actual cost of the R&S policy. In this case, the wall-clock time  $T$  can be used as the cost measure of an R&S procedure. Policies with random run lengths cause  $T$  to be random.

In Sections 2 and 3, the actual run length does not play an important role and the introduced concepts can also be used to evaluate policies with random stopping rules  $\tau$ . Section 4 can be extended to cover policies with random run lengths by considering the expected run length  $\mathbb{E}_{\mathbf{F}}[T]$  (or another functional of the random variable  $T$ ). In this case it is possible to determine whether or not a policy is acceptable to a decision maker, but a different control than  $N$ , such as a confidence parameter, needs to be used to control the policy. For example, in order to implement the popular indifference zone policy  $\mathcal{KN}$  [Kim and Nelson 2001], a confidence parameter  $\eta$ , an indifference zone parameter  $\delta$ , and a first-stage sample size  $n_0$  must be determined. We denote this policy as  $\pi_{\eta, \delta, n_0}^{\mathcal{KN}}$ . Running the  $\mathcal{KN}$  procedure on a given problem will take  $T$  units of wall-clock time, where  $T$  is a random variable. Hence the expected run length of the  $\mathcal{KN}$  procedure depends on  $\eta$ ,  $\delta$ ,  $n_0$  as well as the problem input  $\mathbf{F}$ . For a given scenario  $Q \in \mathcal{Q}$ , the risk is then defined as  $f(\pi_{\eta, \delta, n_0}^{\mathcal{KN}}, Q) = \mathbb{E}_Q[R(\pi_{\eta, \delta, n_0}^{\mathcal{KN}}, \mathbf{F})] - \alpha(Q)$ , and for a given acceptance set  $\mathcal{A}$  it is possible to determine whether  $\pi_{\eta, \delta, n_0}^{\mathcal{KN}}$  is acceptable ( $f(\pi_{\eta, \delta, n_0}^{\mathcal{KN}}, Q) \in \mathcal{A}$ ) or not ( $f(\pi_{\eta, \delta, n_0}^{\mathcal{KN}}, Q) \notin \mathcal{A}$ ). One can further define the best acceptable policy in the class of  $\mathcal{KN}$  policies by optimizing over  $\eta$ ,  $\delta$  and  $n_0$ :

$$\varphi_{\mathcal{A}}(\pi^{\mathcal{KN}}) = \inf_{\eta, \delta, n_0} \left\{ \mathbb{E}_{\mathbf{F}}^{\pi_{\eta, \delta, n_0}^{\mathcal{KN}}} [T] \mid f(\pi_{\eta, \delta, n_0}^{\mathcal{KN}}, \mathcal{Q}) \in \mathcal{A} \right\}.$$

## 5. APPLICATION EXAMPLE

In Waeber et al. [2010] the introduced framework is used to compare three popular policies, namely equal allocation, OCBA [Chen 1996; Chen et al. 2000], and  $(0 - 1)(S)$  [Chick and Inoue 2001b], for the normal R&S problem with  $k = 10$  systems. In this section, we apply the framework to analyze a Bernoulli R&S problem motivated by a concrete simulation optimization application. In particular, we investigate how the preference order on the set of policies  $\Pi$  can change given different assumptions in the performance analysis. Further, we show how the performance measures can be used to identify optimal parameters for an R&S policy.

### 5.1 Allocation of Voting Machines

The simulation optimization problem considered is approximately stated as follows: given a fixed number of touch-screen voting machines, we would like to allocate them to different precincts so that the probability of long waiting times at the polls is small. Often a rough forecast of the turnout is given by the registration of voters. However, factors such as broken voting machines or variation of voters' arrival rates cannot be accurately forecasted. Since the search space grows combinatorially as the number of machines and precincts increases, simulation optimization techniques are helpful in identifying good allocation schemes efficiently. An initial simulation optimization search often yields a set of promising solutions. R&S methods can then be used in a clean-up stage to determine the best system out of the set of promising solutions. This problem, without the clean-up stage, was studied in Allen and Bernshteyn [2006; 2008]. We use this example solely to demonstrate how the performance measures for R&S policies introduced in Sections 2, 3, and 4 can help determine which R&S procedure should be used in the clean-up stage.

Suppose  $k$  promising voting machine allocation schemes  $S_i, i = 1, \dots, k$ , are considered after the simulation search. An allocation scheme is considered good if the probability of long waiting times is small. The performance measure is  $\theta_i = \mathbb{P}_{F_i}(Z_i > C)$  where  $Z_i$  denotes the longest waiting time under  $S_i$  throughout an election day and  $C$  is some quantity, for example, 2 hours. Each allocation scheme produces outputs  $Y_i := \mathbb{1}\{Z_i > C\}$ , so that  $Y_i \sim \text{Bernoulli}(\theta_i)$ .

In many applications there exist options that are not directly associated with a system  $S_i$ , but their performances still depend on the quantities  $\theta_i, i = 1, \dots, k$ . One such option could be “no satisfying system exists and other solutions to the problem need to be identified.” In the voting machine problem, this corresponds to the situation where the probability of long waiting times is so high for all considered allocation schemes that the decision maker needs to consider other undesirable options, such as providing provisional ballots, requesting more voting machines, or petitioning for longer opening hours of the polling stations. We denote this *exit* option as  $i = 0$ , where the performance of  $i = 0$  is a function of the performance of the other  $k$  systems. It is convenient to scale the performance of  $i = 0$  also in  $[0, 1]$ , so we set  $\theta_0 := 1 - \min_{i \in \{1, \dots, k\}} \theta_i$ . Our overriding goal is to identify the option with the *lowest* performance measure, that is,  $\text{argmin}_{i \in \{0, \dots, k\}} \theta_i$ , where only options  $1, \dots, k$  can be simulated. Figure 4 shows the minimal  $\theta_i$  as a function of  $\theta = (\theta_1, \theta_2) \in [0, 1]^2$  when  $k = 2$ . The figure can be interpreted as the performance of an *oracle*, that is, a procedure that always chooses the best option for a given underlying configuration. A good selection procedure should have a performance surface that is close to the oracle performance surface.

### 5.2. Three-Layer Performance Analysis

We can now determine the three layers of the performance analysis for R&S procedures as specified in Section 2.

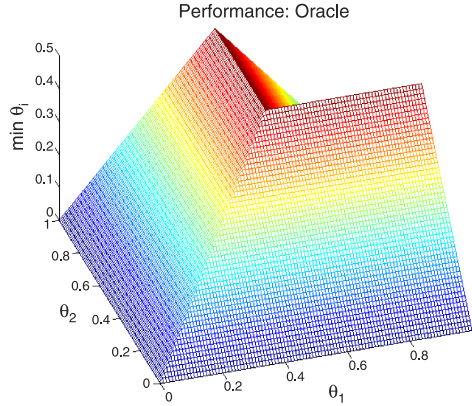


Fig. 4. Performance of an oracle (a procedure that always chooses the best option) for  $k = 2$  systems.

- (1) Use the *linear regret* function as our loss function:

$$L(i_\pi, \theta) := \theta_{i_\pi} - \min_{i \in \{0, \dots, k\}} \theta_i.$$

This loss function reflects the idea that a good R&S procedure minimizes the loss *above* the performance of an oracle procedure.  $L(i_\pi, \mathbf{F})$  corresponds to the *additional* probability of long waiting times attributed to a wrong selection. If this additional probability is small then the selection should still be considered a good selection, even if this not the best available option.

- (2) Use *expected loss* as our configuration-specific risk:

$$R(\pi, \mathbf{F}) := \mathbb{E}_{\mathbf{F}}^\pi [L(i_\pi, \theta)].$$

If the selection procedure  $\pi$  is used for all counties of a state then the expected loss reflects the average loss per county. If the selection procedure is used for a single county then a more risk-averse functional of the distribution of  $L$ , such as a  $q$ -quantile, should be used.

- (3) For the third layer, we consider individually worst-case performance, Bayesian performance, and the new performance measure based on acceptance sets. Different preference orders will be induced by these performance measures. Since we are using linear regret as the loss function, differences between worst-case and indifference zone performances are negligible (as long as the set  $IZ$  is reasonably small).

We compare the following two R&S allocation rules.

- **Equal Allocation.** The fixed simulation budget  $N$  is equally distributed among the  $k$  systems. If  $k$  does not divide  $N$  evenly, the remainder of the simulation budget is allocated to the  $k$  systems according to draws from the  $k$  systems uniformly at random without replacement.
- **2-Stage Allocation.** First, for  $\gamma \in [0, 1]$ ,  $\lceil \gamma N \rceil$  of the fixed simulation budget  $N$  is equally allocated among the  $k$  systems. Second, the  $\lfloor k/2 \rfloor$  systems that performed worst are eliminated and the remaining simulation budget is equally allocated to the rest of the systems. If the allocation in either stage leaves a remainder in the simulation budget, it is assigned according to draws from the set of applicable systems uniformly at random without replacement. In the first stage all  $k$  systems are

applicable, in the second stage only systems that have not been eliminated in the first stage are applicable.

The decision rule is to choose the system with the smallest estimated  $\hat{\theta}_i$  (according to the sample means) as the best system. In case of a tie, one of the systems with the smallest estimated  $\hat{\theta}_i$  is chosen uniformly at random. For the 2-stage procedure, systems eliminated after the first stage cannot be selected as best.

We study two alternatives ( $k = 2$ ) and two policies, which is sufficient to visualize the induced preference orders on  $\Pi$ . The introduced framework can be used to carry out more extensive comparisons (larger  $k$  and more policies), as, for example, is done in Waeber et al. [2010]. In the current article we chose to use a simplified setting and to focus on the structure of the framework and how preference orders on the set of policies can change. Figure 5 shows the estimated expected regret  $R(\pi, \mathbf{F})$  for the equal allocation rule and the 2-stage allocation rule with  $\gamma = 0.7$  as a function of  $\theta_1$  and  $\theta_2$  for different runlengths  $N$ . Note that the expected regret converges pointwise to 0 as  $N \rightarrow \infty$ .

The overall risk, that is, the third layer in the performance analysis process, summarizes each surface displayed in Figure 5 by a single number. Table I reports, for different runlengths  $N$ , estimates of the worst-case performance and two Bayes performances. The first Bayes performance uses a uniform weighting on the plane  $(0, 1]^2$ , while the second uses an independent bivariate weighting with beta-marginals  $\beta(a, b)$  for  $a = 2$  and  $b = 5$ . The density of the  $\beta(a, b)$  distribution is proportional to  $x^{a-1}(1-x)^{b-1}\mathbb{1}_{\{x \in [0, 1]\}}$ , with  $a, b > 0$ . The choice of  $a$  and  $b$  should be set according to the decision maker's risk weighting. We chose  $a = 2$  and  $b = 5$  as an example, and in Section 5.4 we investigate the behavior of the preference order under varying parameters  $a$  and  $b$ .

Explicit calculation of the performance measures is very challenging but estimation results and their corresponding standard deviations appear accurate enough to induce preference orders. The estimation is performed using Monte-Carlo sampling from the  $(0, 1]^2$  plane and stratification with a small stepsize (0.02) is used to reduce variance in the estimation process. 100 uniform samples are drawn from each stratum and a mean performance per stratum is estimated. The worst-case scenario is then estimated as the worst estimated mean performance among all strata. The Bayes performance is estimated by the appropriate weighting of each stratum's estimated mean performance according to the prior risk weighting.

Table I shows that for all estimated runlengths  $N$ ,  $\pi_{2\text{-stage}} > \pi_{\text{equal}}$  for the worst-case analysis, but  $\pi_{\text{equal}} > \pi_{2\text{-stage}}$  for the  $\beta(2, 5)$  Bayes weighting. For the uniform risk weighting the preference order is less obvious. There is some evidence that equal allocation performs better for smaller simulation budgets ( $N = 10, 30$ ), while 2-stage allocation performs better for larger simulation budgets ( $N > 30$ ). This demonstrates that whether a decision maker should use the equal allocation or the 2-stage allocation scheme depends on her risk tolerance and also on the available simulation budget.

### 5.3. Performance Analysis with Acceptance Sets

Similar behavior can be observed while studying different acceptance sets. Figure 6 shows the trajectories of the equal and the 2-stage procedure ( $\gamma = 0.7$ ) under the  $\beta(2, 5)$  Bayes scenario and the worst-case scenario together with three acceptance sets. Table II summarizes the estimated values of  $\varphi_{\mathcal{A}}$  for the two R&S policies under five different acceptance sets. More specifically,  $A_1$  corresponds to a decision maker who focuses on the Bayes  $\beta(2, 5)$  risk weighting;  $A_2$  corresponds to a decision maker who is

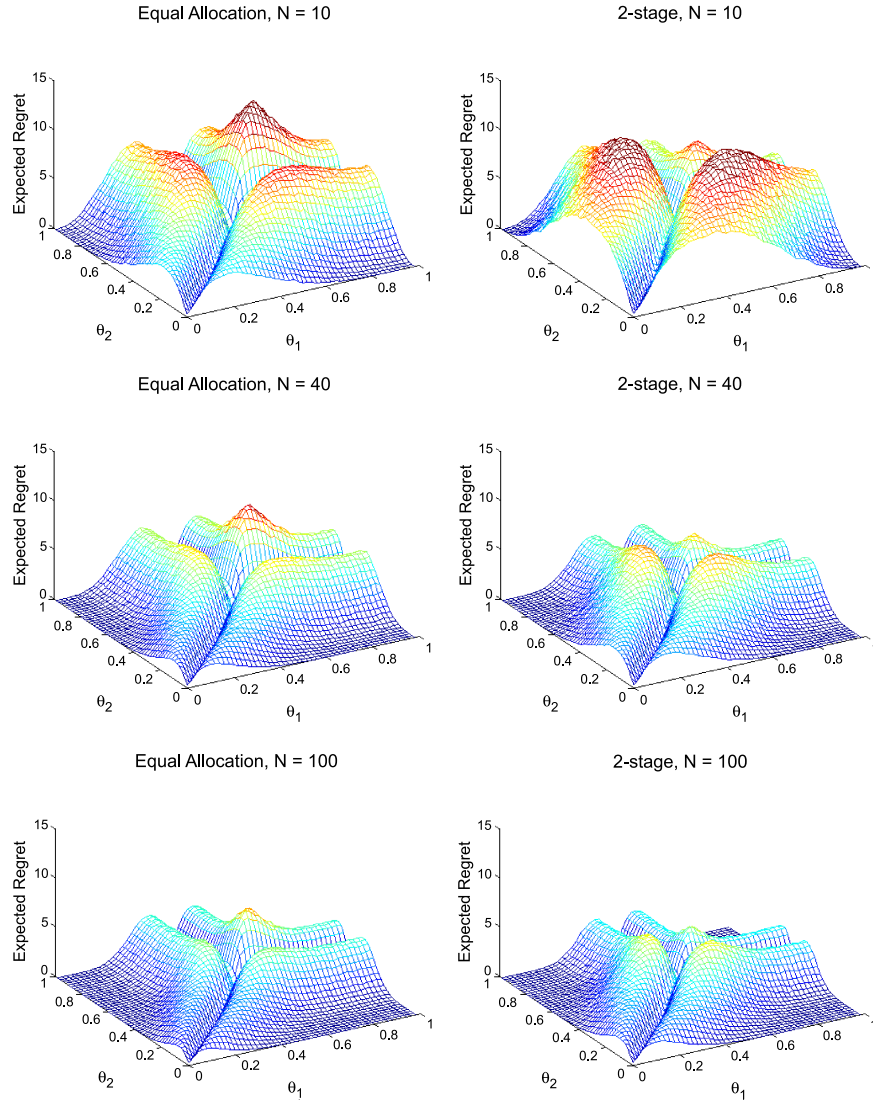


Fig. 5. Estimated expected regret for the equal allocation rule and 2-stage allocation rule with  $\gamma = 0.7$ . As  $N \rightarrow \infty$  the expected regret converges pointwise to 0. The expected regret is given in terms of percentages; to recover absolute values the z-axis should be divided by a factor of 100. The pointwise estimated standard deviation of the expected regret is on average 0.05 and at most 0.16 based on 20,000 iterations for each possible stratum (a square of side 0.02).

concerned about worst-case performance;  $A_3$  corresponds to a decision maker who wants to control both the worst-case and the  $\beta(2, 5)$  risk weighting, but is willing to give up performance in one scenario if the performance in the other is improved;  $A_1 \cap A_2$  describes the most conservative decision maker who requires good performance under both the average-case and the worst-case configuration;  $A_1 \cup A_2$  on the other hand corresponds to a less conservative decision maker who is willing to accept a procedure once its average-case performance or its worst-case performance satisfies some confidence bound.



Table I. Estimated Performance Measures for Different Simulation Budgets  $N$ .

	$N$	10	20	30	40	50	100	Largest std deviation
Worst-Case	Equal	12.34	9.77	7.46	6.69	5.72	4.13	(0.07)
	2-Stage	11.26	7.36	6.11	5.19	4.68	3.27	(0.06)
Uniform risk weighting	Equal	4.71	2.91	1.93	1.54	1.21	0.64	(0.0016)
	2-stage	5.53	2.80	2.00	1.51	1.20	0.61	(0.0018)
$\beta(2, 5)$ risk weighting	Equal	4.28	2.64	1.88	1.49	1.22	0.65	(0.0022)
	2-Stage	6.36	3.39	2.55	1.94	1.63	0.86	(0.0030)

The results are given in terms of percentages; to recover absolute values they should be divided by a factor of 100. The standard deviations are based on a sample of 20,000 iterations. We set  $\gamma = 0.7$  for the 2-stage procedure. These performance measures have been estimated using standard Monte-Carlo methods and stratification to reduce variance.

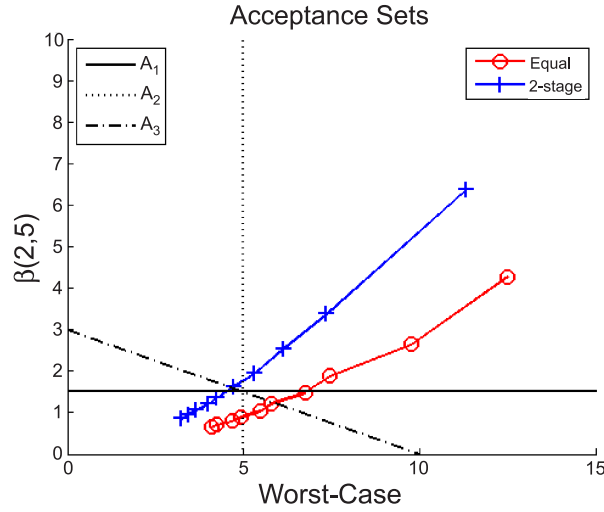


Fig. 6. Trajectories of equal and 2-stage procedures ( $\gamma = 0.7$ ) entering three different acceptance sets. The performance requirement on the Bayesian scenario ( $\rho_{\text{Bayes}}(\pi) \leq 1.5\%$ ) is stricter than the performance requirement on the worst-case scenario ( $\rho_{\text{WC}}(\pi) \leq 5\%$ ). Again, this reflects the idea that a policy with performance in  $A_1 \cap A_2$ , or in  $A_3$ , guarantees good average-case performance *and* its worst-case performance is not extremely bad. See also Figure 3.

Figure 6 and Table II show that a decision maker concerned with the average-case behavior under a  $\beta(2, 5)$  risk weighting would prefer the equal allocation procedure, while a decision maker who focuses on the worst-case performance would choose the 2-stage procedure. This behavior can be explained: the worst-case performance of the equal allocation rule occurs when  $\theta_1 = \theta_2$  and a little bit above 0.5 and hence the exit option would be optimal; see Figure 5. For the 2-stage procedure, on the other hand, the expected regret at these points is much smaller. This makes sense intuitively because the 2-stage procedure focuses in the second stage only on one alternative and the exit option. However, the improvement in worst-case performance comes at a cost: in the region where the exit option is not the best option, the expected regret for the 2-stage procedure is larger than for the equal allocation. This is reflected by the average-case performance for a risk weighting focusing on the region where the exit option is not the best option, such as the  $\beta(2, 5)$  risk weighting.

Table II. Estimated Performance Measures  $\varphi_{\mathcal{A}}$  (along with standard deviations) for Equal and 2-Stage Allocation Rules given Five Different Acceptance Sets

	$A_1$	$A_2$	$A_3$	$A_1 \cup A_2$	$A_1 \cap A_2$
Equal	40 (0.04)	69 (0.79)	49 (0.26)	40 (0.04)	69 (0.79)
2-stage	55 (0.05)	44 (0.79)	51 (0.33)	44 (0.79)	55 (0.06)

We used bootstrapping techniques to estimate  $\varphi_{\mathcal{A}}$  and its standard deviations based on a sample of 200 batches, where each batch consists of an average of 100 simulation runs. The standard deviations are significantly smaller when the entry point into the acceptance set is given by the  $\beta(2, 5)$  scenario. This is due to the fact that the estimation of the worst-case scenario has a much higher variability than the expectation under the  $\beta(2, 5)$  risk weighting.

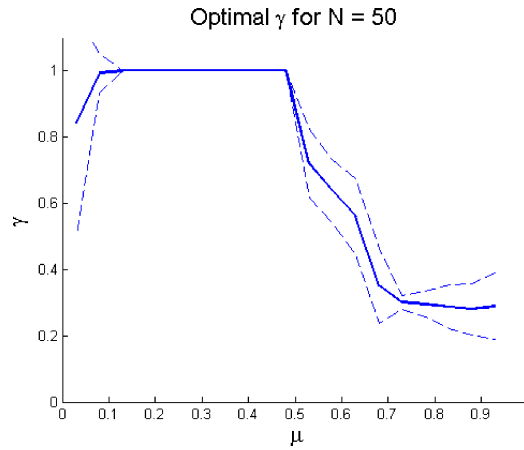


Fig. 7. Optimal selection of the tuning parameter  $\gamma \in [0, 1]$  as a function of the mean  $\mu$  of the  $\beta(a, b)$  risk weighting for a fixed simulation budget  $N = 50$ . The parameters  $a$  and  $b$  are chosen such that the variance is constant and equal to 0.0255, which is possible for parameters  $\mu \in [0.03, 0.97]$ . The optimal parameter  $\gamma$  is estimated by discretization, that is,  $\gamma$  is chosen from the set  $\{i/10, i = 1, \dots, 10\}$  and bootstrapping over a sample of size 20,000. The dashed lines indicate the normal approximation of the 95% confidence interval using the estimated standard deviation from the bootstrap procedure.

#### 5.4 Robustness of Preference Order

We now turn our attention to the robustness of the previous preference orders with respect to the input parameters  $a$  and  $b$  of the beta risk weighting  $\beta(a, b)$ . The mean of the  $\beta(a, b)$  distribution is  $\mu = a/(a + b)$ , and we are interested in how the preference order changes with respect to  $\mu$ . Moreover, for every  $\mu$  an optimal parameter  $\gamma$  of the 2-stage procedure can be found. Figure 7 shows the optimal parameter  $\gamma$  as  $\mu$  varies. The parameters  $a$  and  $b$  of  $\beta(a, b)$  are chosen so that the variance is the same for all choices of  $\mu$ . This constant variance is chosen to be the same variance given by  $a = 2$ ,  $b = 5$ , that is,  $\text{Var} = ab/((a + b)^2(a + b + 1)) = 0.0255$ . As a consequence  $\mu$  can only take values in approximately  $[0.03, 0.97]$ .

The 2-stage procedure with  $\gamma = 1$  corresponds to the equal allocation procedure, and when  $\gamma < 1$  we refer to the resulting 2-stage procedure as a *strict* 2-stage procedure. So for  $\mu < 0.5$ , that is, when  $\mu$  is in the region where the exit option is not the best option, the equal allocation procedure is preferred over a strict 2-stage procedure, but

as soon as  $\mu$  becomes bigger than 0.5, that is, when  $\mu$  is in the region where the exit option is the best option, a strict 2-stage procedure is preferred to the equal allocation procedure. The optimal amount of effort spent in the first stage depends on  $\mu$ . Indeed, the optimal effort  $\gamma$  decreases with  $\mu$ , because a risk weighting with larger  $\mu$  prefers policies that make a correct decision in the region where the exit option is best. Furthermore,  $\gamma$  does not converge to 0 even when  $\mu$  is close to 1. This is because the variance of  $\beta(a, b)$  is kept constant and some risk weighting is still assigned to the region where  $\theta_i < 0.5$ , for  $i = 1, 2$ , that is, where the performance of the policy strongly depends on the outcome of the first stage. Finally, this robustness analysis also depends on the preassigned simulation budget  $N$ . In Figure 7 this is chosen to be  $N = 50$ . For other values of  $N$  the robustness of the optimal parameter  $\gamma$  differs.

## 6. CONCLUSIONS

We presented a performance analysis framework for R&S procedures, which states all assumptions needed to introduce a preference order on a set of R&S policies. The three most popular approaches in the literature have a common representation analogous to convex risk measures used in mathematical finance. Based on this performance evaluation process, we introduced a new performance measure using computational cost and the definition of “acceptable policies.” This easy-to-interpret performance measure can be used to compare selection procedures given different acceptance sets specified by the decision maker. By judicious choice of the acceptance set, a decision maker can identify a performance guarantee that lies between the worst-case and the average-case analysis. Although the framework is presented in the R&S setting where the decision maker wants to select a single best system, the three-layer performance analysis and the concept of acceptance sets could be extended to other simulation optimization settings such as subset selection, constrained and multiobjective optimization, as well as global (continuous) optimization. The key in applying the framework to these settings is to formulate appropriate loss functions  $L$ , configuration-specific risk functionals  $R$ , and performance measures  $\rho$ , but the main structure would likely remain very similar to that presented in this article.

Motivated by a real-world application, we demonstrated how the suggested framework can introduce different preference orders on a set of R&S policies. We further studied the robustness of preference regions under different prior assumptions, which provides us with valuable insight regarding how tuning parameters for an existing R&S policy can be chosen. Such analysis will almost invariably involve heavy computation, as in our example, but can be very helpful in deciding which R&S procedure to use for a given application or class of applications.

## ACKNOWLEDGMENTS

Parts of this work were completed while Shane Henderson was visiting the Mathematical Sciences Institute at Australia National University, and he thanks them for their hospitality and support. The authors would also like to thank conference participants at the 2010 Annual INFORMS meeting and the 2010 Winter Simulation Conference for helpful comments and discussions.

## REFERENCES

- ALLEN, T. T. AND BERNSHTEYN, M. 2006. Optimal voting machine analysis. Tech. rep., DRE Analysis for May 2006 Primary, Steven Hertzberg Ed., Cuyahoga County, Technical Report, Election Science Institute.
- ALLEN, T. T. AND BERNSHTEYN, M. 2008. Helping Franklin County vote in 2008: Waiting lines. Report to the Franklin County Board of Elections.
- ARTZNER, P., DELBAEN, F., EBER, J. M., AND HEATH, D. 1999. Coherent measures of risk. *Math. Finance* 9, 3, 203–228.

- BECHHOFFER, R. E. 1954. A single-sample multiple decision procedure for ranking means of normal populations with known variances. *Ann. Math. Statist.* 25, 1, 16–39.
- BECHHOFFER, R. E., SANTNER, T. J., AND GOLDSMAN, D. M. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. Probability and Statistics. Series, Wiley.
- BERGER, J. O. 1985. *Statistical Decision Theory and Bayesian Analysis*. Springer.
- BOESEL, J., NELSON, B. L., AND KIM, S. H. 2003. Using ranking and selection to “clean up” after simulation optimization. *Oper. Res.* 51, 5, 814–825.
- BRANKE, J., CHICK, S. E., AND SCHMIDT, C. 2007. Selecting a selection procedure. *Manag. Sci.* 53, 12, 1916–1932.
- CHEN, C. H. 1996. A lower bound for the correct subset-selection probability and its application to discrete-event system simulations. *IEEE Trans. Autom. Control* 41, 8, 1227–1231.
- CHEN, C. H. AND LEE, L. H. 2010. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*. World Scientific Publishing Co. Pte. Ltd., Singapore.
- CHEN, C. H., LIN, J., YÜCESAN, E., AND CHICK, S. E. 2000. Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discr. Event Dynam. Syst. Theory Appl.* 10, 3, 251–270.
- CHICK, S. E. 1997. Selecting the best system: A decision-theoretic approach. In *Proceedings of the Winter Simulation Conference*. S. Andradottir, K. J. Healy, D. H. Withers, and B. L. Nelson Eds., IEEE, 326–333.
- CHICK, S. E. AND FRAZIER, P. I. 2009. The conjunction of the knowledge gradient and the economic approach to simulation selection. In *Proceedings of the Winter Simulation Conference*. M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls Eds., IEEE, 528–539.
- CHICK, S. E. AND GANS, N. 2009. Economic analysis of simulation selection problems. *Manag. Sci.* 55, 3, 421–437.
- CHICK, S. E. AND INOUE, K. 2001a. New procedures to select the best simulated system using common random numbers. *Manag. Sci.* 47, 8, 1133–1149.
- CHICK, S. E. AND INOUE, K. 2001b. New two-stage and sequential procedures for selecting the best simulated system. *Oper. Res.* 49, 5, 732–743.
- DEGROOT, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill, New York.
- FÖLLMER, H. AND SCHIED, A. 2002. Convex measures of risk and trading constraints. *Finance Stochast.* 6, 4, 429–447.
- FÖLLMER, H. AND SCHIED, A. 2004. *Stochastic Finance, An Introduction in Discrete Time* 2nd Ed. Walter de Gruyter & Co., Berlin.
- HONG, J. L. AND NELSON, B. L. 2005. The tradeoff between sampling and switching: New sequential procedures for indifference-zone selection. *IIE Trans.* 37, 7, 623–634.
- KIM, S. H. AND NELSON, B. L. 2001. A fully sequential procedure for indifference-zone selection in simulation. *ACM Trans. Model. Comput. Simul.* 11, 3, 251–273.
- KIM, S. H. AND NELSON, B. L. 2006a. On the asymptotic validity of fully sequential selection procedures for steady-state simulation. *Oper. Res.* 54, 3, 475–488.
- KIM, S. H. AND NELSON, B. L. 2006b. Selecting the best system. In *Handbook in Operations Research and Management Science: Simulation*, S. G. Henderson and B. L. Nelson Eds., Elsevier.
- MARKOWITZ, H. M. 1952. Portfolio selection. *J. Finance* 7, 1, 77–91.
- MCNEIL, A. J., FREY, R., AND EMBRECHTS, P. 2005. *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton University Press, Princeton, NJ.
- NELSON, B. L. AND BANERJEE, S. 2001. Selecting a good system: Procedures and inference. *IIE Trans.* 33, 3, 149–166.
- NELSON, B. L. AND MATEJCIK, F. J. 1995. Using common random numbers for indifference-zone selection and multiple comparisons in simulation. *Manag. Sci.* 41, 12, 1935–1945.
- PAULSON, E. 1964. A sequential procedure for selecting the population with the largest mean from  $k$  normal populations. *Ann. Math. Statist.* 35, 1, 174–180.
- RINOTT, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Comm. Statist. Theory Methods* 7, 8, 799–811.
- SANTNER, T. J. AND TAMHANE, A. C. 1984. *Design of Experiments: Ranking and Selection*. Marcel Dekker, Inc.
- SOBEL, M. AND HUYETT, M. J. 1957. Selecting the best one of several binomial populations. *Bell Syst. Techn. J.* 36, 537–576.

- SWISHER, J. R., JACOBSON, S. H., AND YÜCESAN, E. 2003. Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: A survey. *ACM Trans. Model. Comput. Simul.* 13, 2, 134–154.
- WAEBER, R., FRAZIER, P. I., AND HENDERSON, S. G. 2010. Performance measures for ranking and selection procedures. In *Proceedings of the Winter Simulation Conference*. B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan Eds., IEEE, 1235–1245.

Received April 2011; revised March 2012; accepted May 2012